

## SYSTEM TO MAP REMOTE LAN HOSTS TO LOCAL IP ADDRESSES

This invention relates to a system for mapping remote hosts located on remote local area networks connected to the internet to appear as if they are devices attached to a single local area network (LAN).

### **Background of the invention**

The protocol used for internet communications is TCP/IP. TCP/IP is actually a suite of protocols that work in conjunction with one another to provide complete communications services between and across networks. When a data packet is to be delivered from a host device on a local network to a host on a remote network, the packet must be “addressed” to the logical address of the destination device, and must then be routed from the local network to that destination device through a series of jumps, or “hops” across network segments connected by routers until it reaches the destination network, where it is sent to the destination device. “IP addressing” is used to identify the destination network and receiving host on that network.

Since its inception, TCP/IP (IPv4) has used a 32-bit addressing scheme (the “IP address”) to uniquely identify devices attached to the internet. These addresses, consisting of 32 bits grouped into 4 octets, are easily recognized in their familiar xxx.xxx.xxx.xxx decimal format, which is well-known to users of the internet.

Each IP address includes a network identification (netID) and a host identification (hostID). The netID may comprise between the first 8 bits to 24 bits of the IP address, with the remainder indicating the hostID, depending upon whether the network is a class A, B, or C network. In class A networks, the first octet is the net ID, and the last three are the hostID. For class B networks, the netID is the first two octets, and the hostID is the last two octets. Class C networks have a netID

of the first three octets, with the last octet being the hostID. An IP network consists of a group of hosts that share the same netID. Through this addressing scheme, groups of contiguous IP addresses may be segregated into IP networks as LANs, and may be connected to each other or to the internet by routers.

Routers operate at the network and datalink layers of the OSI (“open systems interconnection”) model, and are used to route IP packets (also called “datagrams”) between IP networks, and to deliver or receive them to and from a host on the same network segment. Communications within a single IP network segment take place at the data link layer, and use physical addresses, also hardware addresses, to deliver the packet to the correct host on the segment.

IP Packets formatted for transmission across a network are carried across the network within a “frame” whose header has fields for the source and destination physical addresses of the sending and receiving hosts. Each type of network (ethernet, token ring, frame relay, etc.) has its own frame type. On ethernet networks, physical addresses may also referred to as “media access control,” or “MAC” addresses. Each network communicating device, such as a network interface card, or “NIC,” has a physical address encoded within it by the manufacturer, uniquely identifying the manufacturer and the device. Thus, when an IP packet is encased within a frame, the frame is used only to transport the packet between hardware addresses on a single network. When the IP packet is to be sent from one network to another, the frame is received at a router connecting the two networks, is broken down to determine the ultimate IP address. The IP packet is then encased within another frame for delivery across the second network to the next device, which may be another router or an ultimate destination.

When an IP packet is formatted for delivery, the sending host will first check the IP address to see whether the destination host is located on the same network. If it is not, the sending host will

1 send the packet to a default router by placing the router's hardware address in the header of the  
2 frame and placing the packet on the local network. The router will receive the packet and, using one  
3 of a number of various routing protocols, determine the next router along the path that the packet  
4 will be delivered to. This process is repeated as a series of hops until the packet arrives at a router  
5 connected to the destination network.

6 If the sending host determines that the packet is destined for a host on the same network, the  
7 sending host will send the packet directly to the physical address of the receiving host using address  
8 resolution protocol, or "ARP," to determine the correct physical address. It does this by first  
9 consulting its cache of recently resolved physical addresses to see whether the hostID part of the IP  
10 address is cross referenced to a physical address. If the physical address is found within the cache,  
11 it is then written to the destination field of the frame header and sent to the network. If the physical  
12 address is not found within the cache, the sending host will broadcast an ARP request to all devices  
13 on the network to determine the physical address of the device to which the IP packet is addressed.  
14 Upon receiving an ARP request, the device having that IP address will reply by sending its hardware  
15 address to the requesting host. All other devices will ignore the ARP request. Upon receiving the  
16 reply from the destination host, the sending host will place the physical address in the destination  
17 field of the frame header and send the packet to the network.

18 Because of these routing and address properties and limitations, the following guidelines  
19 should be followed when setting up or attaching to a network: All hosts having the same netID  
20 should be attached to the same network segment so that they can communicate directly. Hosts  
21 separated by a router will be unable to communicate even though they have the same netID. Hosts  
22 with different netIDs *must* communicate through a router, even though they may be attached to the  
23 same network segment. As used in this description, an "IP network" shall refer to a contiguous

1 network in which all actual and virtual devices have the same netID and are not required to  
2 communicate through a router.

3 Internet communications take place through the sending and receipt of internet protocol (IP)  
4 “packets” across the communications medium. Each IP packet must have an IP header, which is a  
5 block of information providing, among other things, a destination IP address for the packet and the  
6 source IP address of the sending device. By analyzing this information, internet routers are able to  
7 deliver packets to the networks indicated by the IP address.

8 Every device connected directly to the internet must be assigned a globally unique IP address  
9 in order that communications between the device and others on the internet may take place.  
10 However, with the explosive growth of the internet in the 1990's, the number of available addresses  
11 for new devices became critically low. One solution to the problem of limited global addresses was  
12 to cease assigning a global IP address to every device on a LAN connected to the internet. Rather,  
13 through the development of network address translation technology, a LAN can be isolated from the  
14 internet by a network address translator (NAT) that is attached both to the LAN and to the internet.  
15 Using a NAT, hosts on the LAN can be assigned private IP addresses that are unique within the  
16 LAN, even though they are not globally unique. Such addresses may also be used by any other  
17 LAN, whether or not it is connected to the internet. However, any LAN using private IP addresses  
18 must also have a NAT if it is attached to the internet.

19 In operation, the NAT is assigned a global IP address that enables it to be seen from the  
20 internet, and is also assigned a local IP address that can be seen by other devices on the LAN. It is  
21 common for a NAT to be assigned more than one global IP address, in which case the NAT will  
22 have a pool of IP addresses from which to select one that is available. When a local host wishes to  
23 communicate across the internet, it sends an IP packet having in its destination field the global IP

1 address of the intended recipient, and placing its own local IP address in the source field. As the  
2 packet passes through the NAT connecting the LAN to the internet, the NAT substitutes  
3 (“translates”) its own global IP address in the source field of the IP packet. Following the address  
4 translation, the NAT will forward the packet to the internet for routing and delivery. At the same  
5 time, the NAT saves information in its internal tables sufficient to enable it to identify a response  
6 from the intended recipient. When a response having as a destination address the source address  
7 substituted by the NAT is received at the NAT, the NAT will analyze the packet to ensure that it  
8 represents a valid reply, and if it is valid will “translate” the packet’s destination address to the local  
9 host’s IP address and forward the packet to the LAN. A more complete description of the operation  
10 of a NAT may be found in U.S. Patent No. 5,793,763 to Mayes et al., (“Security System for Network  
11 Address Translation Systems”).

12 In the internet IP addressing scheme, three blocks of contiguous IP addresses have been  
13 withdrawn from “public” use as global IP addresses, and are reserved for “private” use on LANs  
14 isolated from the internet. The reserved private address ranges are:

15 10.0.0.0 to 10.255.255.255 (The network 10.xxx.xxx.xxx is a single class A network  
16 supporting up to 16,777,214 host addresses);

17 172.16.0.0 to 172.31.255.255 (The networks 172.16.xxx.xxx to 172.31.xxx.xxx are 16 class  
18 B networks, each supporting up to 65,534 host addresses); and

19 192.168.0.0 to 192.168.255.255 (The networks 192.168.0.xxx to 192.168.255.xxx are 256  
20 class C networks, each supporting up to 254 host addresses).

21 Using these address spaces, any isolated LAN can utilize a contiguous block of private  
22 addresses that is large enough to fit its needs. These private address spaces provide more than  
23 enough addresses to accommodate even the largest enterprise LANs.

1 NAT not only makes it possible for computers on an isolated LAN using private IP addresses  
2 to communicate with other computers on the internet, but it is also a security measure for isolating  
3 devices on the LAN from the internet and for keeping unwanted transmissions from reaching  
4 devices on the LAN. Most NATs are programmed to reject “unexpected” packets sent to their  
5 LANs. Such packets could include, for example, transmissions attempting to initiate a TCP session  
6 from outside the LAN, UDP packets that are not in reply to a domain name service (DNS) request  
7 initiated by a local host, certain internet control message protocol (ICMP) packets, and network file  
8 system (NFS) packets, which are designed to access an external computer’s file system as if it were  
9 local.

10 Even though NATs are intended to protect their LANs from potentially harmful external  
11 attacks, there are a number of “good” reasons why a LAN administrator might wish to allow certain  
12 externally initiated transmissions into the LAN. At least one of these reasons is to permit  
13 monitoring and maintenance of devices on a remote LAN by a computer consultant or network  
14 administrator. Other circumstances in which it may be desirable to allow externally initiated  
15 sessions to pass through to the LAN may include external auditing of a business or records by an  
16 authorized accounting firm and inventory control of remote business sites from a single  
17 administrative location. However, the very services that such uses would require to monitor  
18 computer and records maintained on a remote LAN are precisely those uses that NATs are designed  
19 to reject. One method for overcoming this obstacle is to use a tunneling protocol between the  
20 remote LAN and the local NAT. Where the initiating and receiving devices are on separate LANs,  
21 tunneling protocols can be established between two trusted NATs to allow transmissions from a host  
22 on one LAN to be received by a host on a second, using the internet as the transmission medium.  
23 This type of communication may also incorporate encryption and authentication security measures

1 to protect the integrity of the transmissions.

2 Tunneling is a process in which a packet being transmitted between remote hosts may be  
3 encapsulated as a payload within another packet for transmission between two trusted gateways or  
4 other endpoints of the tunnel. An original packet is sent from the originating host to the trusted  
5 device, where it is enclosed as the payload of a new IP packet, and a new IP header is prepended to  
6 it with its destination field containing the IP address of the device at the end of the tunnel. Upon  
7 arrival at the end of the tunnel, the new “outer” header is stripped away, and the original packet may  
8 then be forwarded to a LAN or further processed, as appropriate. By using a tunnel, it is possible  
9 to circumvent conventional routing mechanisms for the encapsulated packet during transit, while  
10 it is in the tunnel.

11 These prior art devices and methods are useful in making resources located on the internet  
12 and on remote LANs available to other devices. However, while the methods for accessing remote  
13 hosts are available, they tend to be awkward and time consuming to use because of the requirement  
14 for recording, saving, and re-entering the IP addresses of such hosts whenever access is desired.  
15 Where many hosts located on separate and independent LANs are to be accessed from a single  
16 “home” location, it is necessary for each remote device to be separately accessed across the internet.  
17 This is not only inefficient, but promotes the likelihood of error whenever a remote IP address must  
18 be manually entered. In addition, when two or more remote LANs are using identical IP addressing  
19 schemes, a problem arises when trying to route IP packets to those LANs since they are not uniquely  
20 addressable. What is needed is a network design in which local IP addresses on remote LANs may  
21 be translated into addresses that are unique, and thus routable, as seen from a “home” LAN. With  
22 such a design, remote hosts on isolated LANs can be seen and accessed from a “home” computer  
23 as if each remote host were situated upon the “home” computer’s LAN, and the “home” LAN has

1 a structured design such that remote LANs appear to be logical, contiguous subnets that are part of  
2 a single, manageable network.

### 3 **Summary of the invention**

4 The present invention uses routers configured to operate as virtual-private-network routers,  
5 or VPN-routers, to connect computers on a “home” LAN to hosts on any number of separate and  
6 isolated LANs such that each remote host appears to be a local computer connected to the “home”  
7 LAN. VPN-routers combine the functions of network address translation, encryption and  
8 authentication, routing, tunneling, and address resolution protocol to determine physical addresses.  
9 In addition, VPN-routers must be able to perform normal routing and network address translation  
10 for communications with other devices on the internet that are not part of the internetwork  
11 addressing scheme of this invention.

12 In this invention, IP tunneling is used in conjunction with network address translation to  
13 create a tunnel from the “home” LAN’s VPN-router to the VPN-router for each isolated LAN having  
14 hosts that will be accessed. Each remote LAN uses local IP addresses taken from the blocks of  
15 reserved private IP addresses, although this is not strictly a requirement of the invention. The  
16 “home” VPN-router and the remote VPN-router may be preconfigured to incorporate encryption and  
17 authentication security. Because they are preconfigured, it will not be necessary for them to  
18 negotiate security keys in the clear each time a new session is opened.

19 The present invention utilizes the internal tables of the VPN-routers for the “home” LAN and  
20 each remote LAN that is appear as a local network to perform the required address translation to  
21 ensure proper delivery and receipt of “home”-LAN-to-remote-LAN communications. Because the  
22 translation tables for actual and virtual IP addresses used in the addressing scheme of this invention  
23 are preconfigured, each attached device included in the scheme must be assigned a static IP address,

1 rather than having an address automatically assigned using dynamic host control protocol (DHCP).  
2 Devices attached to a LAN that are not part of the design of this invention may use DHCP without  
3 adversely impacting upon hosts that are included in the design of this invention. In one  
4 embodiment, each remote VPN-router will constitute one end of a single tunnel, the other end of  
5 which will be the “home” VPN-router. The “home” VPN-router will comprise one end of as many  
6 tunnels as there are remote LANs. Other embodiments may use the system of this invention to  
7 provide multiple connections between remote LANs, although the complexity of such a system will  
8 increase as the number of tunnels connecting remote LANs increases.

9 Each LAN is provided with an addressing scheme in which each local and remote host that  
10 will appear on the LAN is given a locally unique private IP address. Addresses assigned to hosts  
11 actually connected to a LAN will be referred to as “actual” addresses, while addresses assigned to  
12 remote hosts will be referred to as “virtual” addresses. In one embodiment, the system administrator  
13 may assign actual and virtual IP addresses from a block of contiguous private IP addresses that is  
14 large enough for each host to have a unique address. In other embodiments, particularly those in  
15 which legacy network addresses have been established and will be used without further  
16 modification, virtual addresses must be taken from the unused address space within those local  
17 addressing schemes.

18 Once the address schemes for all LANs within the system have been determined, each VPN-  
19 router must have its internal tables preconfigured to hold all IP addresses to be used on its LAN in  
20 the addressing design of this invention. It is possible that some actual hosts on a remote LAN will  
21 not be included in the design of this invention, and as to those hosts, it will not be necessary for their  
22 IP addresses to be preconfigured in the VPN-router: if desired, they may use DHCP. Each actual  
23 host on a LAN must have its actual address cross referenced with that host’s virtual address on the

1 remote LAN (the LAN at the other end of the tunnel), and each virtual address on the LAN must be  
2 cross referenced to the actual address of the host on the remote LAN. The "home" VPN-router must  
3 have a set of internal tables corresponding to each of the remote LANs to which it is connected. For  
4 each remote LAN, the "home" VPN-router will maintain a table cross referencing the addresses of  
5 actual hosts on the "home" LAN with the virtual addresses of those hosts on the remote LAN, and  
6 cross referencing the virtual addresses of remote hosts with the actual addresses used by those hosts  
7 on the remote LAN. Where there are two or more "home" LANs, the addressing scheme will simply  
8 be applied to each of them. In such a design, each "home" LAN will be treated as a remote LAN  
9 by every other "home" LAN.

10 Once the IP addressing scheme has been created and implemented, each remote host will  
11 appear as a virtual host having a unique virtual IP address on the "home" LAN, and each host on the  
12 "home" LAN will appear as a virtual host having a unique virtual IP address on each remote LAN.  
13 However, hosts on one remote LAN will normally be isolated from other remote LANs, and will not  
14 be able to see hosts on other remote LANs unless the network is specifically designed to provide  
15 such visibility. In that event, a remote LAN will become a "home" LAN with respect to any other  
16 remote LANs it has been configured to "see."

17 When an IP packet is to be sent from a host on the "home" LAN to a host on a remote LAN,  
18 the packet is first delivered to the "home" VPN-router, using the hardware address for that router.  
19 Because the VPN-router "understands" that it may be working with virtual IP addresses, it will  
20 respond to ARP requests directed to a virtual host by providing its own hardware address. Upon  
21 arriving at the VPN-router, the packet is analyzed and is determined to be either a "normal" packet  
22 destined for normal routing and delivery across the internet, or is a "special" packet subject to the  
23 method of this invention. If it is a "normal" packet, it will undergo standard processing, including

1 network address translation, and will be forwarded to the internet for delivery. If it is a “special”  
2 packet, it’s actual source address will be translated to be the unique, private virtual IP address that  
3 references the “home” host within the remote LAN’s addressing scheme. The packet will then be  
4 then encrypted (if desired), and encapsulated within an outer IP packet in which the destination field  
5 contains the global IP address of the remote VPN-router, and the source field contains the global IP  
6 address of the local VPN-router. The packet will then be then forwarded to the internet, which will  
7 route it to the remote NAT gateway. Upon being received by the remote VPN-router, the packet’s  
8 outer IP header will be stripped, and the destination and source IP addresses are analyzed. The  
9 remote VPN-router will find that the destination address is a virtual, and not an actual address of a  
10 host on its LAN, and will then translate the virtual destination address to the actual address of that  
11 host on its LAN.

12 A responding packet from the remote host will have as its destination address the virtual  
13 address of the “home” host, and will include its own actual address in the source field. When the  
14 packet reaches the remote VPN-router, the actual source address will be replaced by the virtual  
15 source address of the host as seen by the “home” LAN, and the packet will be encrypted (if desired),  
16 encapsulated, and sent along the tunnel to the “home” VPN-router. Upon arrival there, the packet  
17 will be decapsulated, and the “home” VPN-router will recognize the destination as being a virtual  
18 destination for an actual device on the “home” LAN, and will translate the virtual destination  
19 address to the actual address before sending the packet across the “home” LAN.

## 21 **Brief Description of the Drawings**

22 Figure 1 depicts an internetwork in which a “home” LAN is connected to three remote LANs  
23 across the internet.

Figure 2 depicts the internetwork of Figure 1 from the perspective of the “home” LAN using the method of this invention.

Figure 3 shows the internetwork of Figure 1 from the perspective of one of the remote LANs.

Figure 4 depicts an example Network Address Scheme for the method of this invention.

Figure 5 provides an example of VPN-router translation tables for each of the VPN-routers in the network shown in Figure 1.

Figure 6 illustrates example address translations for hypothetical transmissions in accordance with this invention.

Figure 7 illustrates the decision tree used by a VPN-router to route IP packets from a LAN in accordance with the method of this invention.

Figure 8 illustrates the decision tree used by a VPN-router to route IP packets from the internet in accordance with the method of this invention.

## Detailed Description of the Preferred Embodiments

Figure 1 illustrates a typical network configuration in which a “home” LAN 10 (which will also be referred to as “Network A”) is connected to the internet 60. Remote LANs 70 (“Network B”), 120 (“Network C”) and 170 (“Network D”) are also connected to the internet, but are not connected to each other except across the internet. According to the method of this invention, Networks “A,” “B,” “C,” and “D” comprise an intranetwork that may be regarded as a virtual LAN on a single network segment, as viewed from the perspective of Network “A.” As viewed from the perspective of any remote network, the LAN will appear to include only that remote network and any “home” network. In Figure 1, Network “A”'s VPN-router 50 connects hosts 20, 30 and 40 to the internet, but also acts as a firewall to isolate them from the internet and any other network from

1 receiving unwanted communications. Network "A" is the network with which all other networks  
2 must be able to communicate as if attached to a LAN, and the local addressing scheme for Network  
3 "A" reflects an overall design that is expandable as necessary to incorporate virtually any number  
4 of additional networks into the internetwork.

5 NAT-routers 80, 130, 180 connect Networks "B," "C," and "D," respectively, to the internet  
6 in the same manner. Each VPN-router has a global IP address at its interface with the internet, and  
7 a local IP address at its interface with its LAN. Network "B" 70 has three hosts, 90, 100 and 110,  
8 identified in Figure 1 by their local IP addresses. Network "B" is a legacy network which has  
9 retained the private IP addresses that were assigned prior to the incorporation of Network "B" into  
10 the intranetwork. Network "C" has three hosts, 140, 150 and 160, each being identified by its local  
11 IP address, and also having legacy local IP addresses. Hosts 140 and 150 on Network "C" have  
12 local IP addresses identical to those of hosts 90 and 110 on Network "B." However, in accordance  
13 with the method of this invention, there is no ambiguity when these hosts are included in a virtual  
14 LAN, as described below. Network "D" has three hosts, 190, 200 and 210, which also are identified  
15 by their local IP addresses. Network "D," however, uses local IP addresses that were assigned in  
16 accordance with an overall intranetwork naming scheme. As a result, it will not be necessary for  
17 hosts on Network "D" to have their addresses translated for communications with Network "A."

18 When configured in accordance with the method of this invention, the intranetwork of Figure  
19 1 will take on the virtual topography illustrated in Figure 2, as seen from Network "A." The virtual  
20 LAN of Figure 2 follows the addressing scheme shown in Figure 3. In accordance with this scheme,  
21 hosts on Network "A" may be assigned local IP addresses in the range of 10.200.1.1 up to  
22 10.200.1.255. Because this network is a class A network having a NetID = 10, there is no  
23 prohibition against assigning an address having a fourth octet of "255" so long as the second and

third octets are not also “255,” in which case the address would be reserved for broadcasts across the network. Network “B” will be assigned virtual IP addresses in the range of 10.200.2.1 up to 10.200.2.255. Again, because these addresses are on the class A network having a NetID = 10, they will be directly accessible from other hosts on the network. The use of a “2” in the third octet of the addresses is for administrative convenience and ease of reference. Similarly, Network “C” uses virtual IP addresses ranging from 10.200.3.1 up to 10.200.3.255, while Network “D” uses virtual IP addresses from 10.200.4.1 up to 10.200.4.255. As appears in Figure 2, all hosts on the intranetwork appear as hosts attached to a single segment LAN, making it possible for the to communicate directly with one another using their physical addresses. From the perspective of Network “A,” VPN-router 50 provides a connection to the internet for all communications other than those destined for actual or virtual hosts on the LAN. However, according to the method of this invention, communications between hosts on Network “A” and other actual or virtual hosts on the LAN will take place as if the hosts were all attached to the same network segment.

Figure 3 depicts the way the internetwork of Figure 1 would appear from the perspective of Network “B.” Since Network “B” uses local IP addresses that are different from the local IP addresses used by Network “A,” all actual and virtual hosts on the “B” Network will have the local IP addresses assigned for that network.

Figure 4 shows the overall network address scheme. Actual hosts on Network “A” 10 have the actual IP addresses listed for that network. Actual IP addresses for Networks “B” 70, “C” 120, and “D” 170 are listed under “Actual LAN IP Addresses.” Virtual addresses for those networks and the hosts attached to them, as seen from Network “A” are listed under “Virtual LAN IP Addresses on Network A.” Virtual IP addresses of the hosts on Network “A,” as seen locally from the other networks are listed under “Virtual LAN IP Addresses on Local LAN.” The global IP addresses of

1 the VPN-routers are listed under "Internet IP Addresses (Global)." Each host on each network has  
2 been designated by a number (Host 1, Host 2, etc. . .) for ease of reference. The host numbers,  
3 however, are simply illustrative references, and have nothing to do with the addressing scheme of  
4 this invention.

5 As shown in Figure 4, each host on Networks "B," "C," and "D" has been assigned a virtual  
6 IP address by which it can be referenced from Network "A." Figure 4 also shows that the virtual  
7 addresses assigned to Network "D" are the same as the actual local IP addresses for that network.  
8 As a general rule, where remote networks are to be incorporated into the internetwork design of this  
9 invention, local IP addresses should be assigned to correspond with the virtual IP addresses for that  
10 network unless other considerations (such as a desire to maintain an earlier addressing scheme, or  
11 the need to maintain compatibility with other parts of a pre-existing LAN) outweigh that choice.

12 Because hosts on Networks "B," "C," and "D" must be able to send packets to hosts on  
13 Network "A," virtual IP addresses must be assigned to those hosts from the available address space  
14 for each of those networks. In accordance with this requirement, virtual IP addresses have been  
15 assigned to hosts 1, 2 and 3 from unused addresses on Network "B": Host 1 has the virtual IP  
16 address "192.168.10.10"; host 2 has the virtual IP address "192.168.10.11"; and the virtual IP  
17 address "192.168.10.12" is assigned to host 3. A similar scheme is used to assign virtual IP  
18 addresses to hosts 1, 2 and 3, as seen from Network "C." However, because Network "D" was  
19 designed from the ground up to fit within the addressing scheme for the intranet of this invention,  
20 the actual addresses for hosts 1, 2 and 3 on Network "A" can also serve as the virtual addresses for  
21 those hosts, as seen from Network "D."

22 Internal translation tables for the VPN-routers shown in Figure 1 are given in Figure 5.  
23 Because Network "A" must be able to communicate directly with Networks "B," "C," and "D," it

1 must have a translation table for packets destined to each of those networks. The purpose for  
2 translation by VPN-router "A" is to replace the source address in the packet's IP header with the  
3 virtual address of the host on Network "A" from which the packet originated. In this manner, a  
4 reply packet from the remote network will be able to use the source address from the packet it  
5 received as the destination address for the reply packet it will send. The packet will then be  
6 encapsulated and sent VPN-router to VPN-router via the internet, where it will be routed according  
7 to standard routing protocols. Upon arrival at the receiving VPN-router, the packet will be  
8 decapsulated, and the "inner" IP packet will have its destination address translated in accordance  
9 with the receiving VPN-router's translation table. As shown in Figure 5, the VPN-routers for  
10 Networks "B," "C," and "D" have only one set of translation tables each, which will be sufficient  
11 for the internetwork shown in Figure 1. However, if it is desired that any of those networks should  
12 be able to communicate directly with any other one of them, a second set of translation tables would  
13 have to be added to the VPN-router for each to perform the necessary address translation. In this  
14 event, additional unused local IP addresses would have to be available for assignation to the virtual  
15 hosts to be added to each network.

16 Figure 6 provides eight examples of the method of address translation of this invention. At  
17 220, a packet is sent from Host 1 on Network "A" to Host 5 on Network "B." As the packet leaves  
18 Host 1, its source field contains the actual IP address of Host 1, and its destination field contains the  
19 virtual IP address of Host 5, as seen from Network "A." When the packet reaches the sending VPN-  
20 router, which in this case is VPN-router "A," its source address is translated to Host 1's virtual IP  
21 address, as seen from Network "B" ("192.168.10.10"). VPN-router "A" may encrypt the packet,  
22 or provide authentication information for security, and will then encapsulate the packet for transit  
23 along the tunnel to VPN-router "B." The encapsulation header will also be an IP header, and the

1 source and destination addresses will be the global IP addresses of the two VPN-routers that are the  
2 endpoints of the tunnel, as shown in the column headed "Tunnel Routing." At the receiving VPN-  
3 router (VPN-router "B"), the packet is decapsulated and, if necessary, decrypted and authenticated.  
4 VPN-router "B" will then translate the destination address to be the actual IP address of Host 5, in  
5 accordance with its translations tables, and will send the packet to Network "B." (As previously  
6 described, VPN-router "B" will actually send the packet to the physical address of Host 5 although,  
7 for purposes of this invention, this step will be transparent). The packet will be received by Host  
8 5, and will bear the actual IP address of that host in its destination field and the virtual IP address  
9 of Host 1 in its source field.

10 At 230, Host 5 sends a responsive packet back to Host 1. The same procedure is followed  
11 in which the packet will have the actual IP address of Host 5 in its source field, and the virtual IP  
12 address of Host 1, as seen from Network "B," in its destination field. At VPN-router "B," the source  
13 field of the packet is translated to be the virtual address of Host 5, as seen from Network "A," and  
14 the packet is processed for security, encapsulated, and routed to VPN-router "A." At VPN-router  
15 "A" the packet is decapsulated, processed for security, and the destination IP address is translated  
16 to be Host 1's actual IP address. Similar address translation occurs for transmissions 240 through  
17 290. It may be noted that transmissions at 240 - 250 involve Host 7 on Network "C" while  
18 transmissions at 280 - 290 involve Host 4 on Network "B." Although these hosts have identical  
19 actual IP addresses on their respective networks, there is no conflict or ambiguity in the sending or  
20 receipt of transmissions to these hosts because each has a unique virtual IP address as seen from  
21 Network "A."

22 The VPN-router of this invention must be specially configured to hold virtual IP addresses  
23 in its translation tables, and to identify IP packets being transmitted to or from the virtual hosts

1 having virtual IP addresses. This behavior is different from the standard behavior of a NAT or a  
2 router, and must be specifically designed into the VPN-router of this invention. All VPN-routers  
3 used to isolate the networks comprising the internetwork of this invention may be the same in their  
4 hardware and firmware, and would differ only in the mapping of actual and virtual IP addresses in  
5 accordance with the addressing scheme of this invention.

6 Figure 7 depicts a decision tree that a VPN-router would use to route IP packets from its  
7 local network. The process begins at 300, and at 310 a packet is received from the local network.  
8 The packet is first examined 320 to see whether it is destined for an IP address appearing on the  
9 local network, or should be sent to the internet. If it is destined for delivery to the internet, "normal"  
10 source address translation 330 and security measures 340 will be applied and the packet will be  
11 delivered to the internet for further routing. However, if the packet is destined for the local network,  
12 the NAT router will then consult its internal tables to determine whether the destination IP address  
13 is actual or virtual 350. If the packet is being sent to an actual IP address on the local network, the  
14 receiving host will be actually attached to the network, and will receive the packet without any  
15 action being taken by the VPN-router, which may ignore the packet 360. If the destination address  
16 is a virtual address, the VPN-router must determine to which local network in the intranetwork the  
17 packet should be routed 370. In the example intranetwork illustrated in Figure 1, if the VPN-router  
18 is attached to Networks "B," "C," or "D," then the only virtual network to which the packet could  
19 be routed is Network "A." If, however, the packet is received by VPN-router "A" from Network  
20 "A," then it will be necessary for the router to determine whether the virtual IP address is on  
21 Network "B," "C," or "D." A similar determination will also need to be made for intranetworks in  
22 which more than one network is configured as a "home" network.

23 Once the network to which the packet will be forwarded has been identified, the VPN-router

1 will substitute the virtual IP address of the sending host 380. If encryption 390 (or other security  
2 measures) have been activated, the VPN-router will perform the encryption 400 or other security  
3 measures. The VPN-router will next encapsulate the packet within an IP packet addressed to the  
4 VPN-router for the destination network 410, and will deliver the encapsulated packet to the internet  
5 420 for routing to the destination VPN-router.

6 Figure 8 illustrates the decision tree for IP packets received from the internet by a VPN-  
7 router. The process begins at 430, and at 440 a new packet is received from the internet. The packet  
8 must be decapsulated 450 and inspected. If it is encrypted 460, it must first be decrypted 470 before  
9 further analysis can take place. If the IP packet, now free of encapsulation, has an actual destination  
10 IP address, then the packet will be processed “normally,” that is, it will be authenticated 490 and  
11 destination address translation 500 will be performed. Here, it may be noted that, in some  
12 implementations the actual destination IP address for a “normal” packet may be the VPN-router’s  
13 global IP address, which will then be translated in accordance with other criteria maintained by the  
14 VPN-router. However, where the decapsulated packet has a virtual IP address in its destination  
15 field, the VPN-router will translate that to be the receiving host’s actual local IP address, and will  
16 deliver the packet to the local network for delivery to the host.

17 While the invention has been described, disclosed, illustrated and shown in various terms  
18 or certain exemplary embodiments or modifications which it has assumed in practice, the scope of  
19 the invention is not intended to be, nor should it be deemed to be, limited thereby and such other  
20 modifications or embodiments as may be suggested by the teachings herein are particularly reserved  
21 especially as they fall within the breadth and scope of the claims here appended.  
22